

# Gradual Release of Sensitive Data under Differential Privacy

Fragkiskos Koufogiannis, Shuo Han, George J. Pappas\*

April 3, 2015

## Abstract

We introduce the problem of releasing sensitive data under differential privacy when the privacy level is subject to change over time. Existing work assumes that privacy level is determined by the system designer as a fixed value before sensitive data is released. For certain applications, however, users may wish to relax the privacy level for subsequent releases of the same data after either a re-evaluation of the privacy concerns or the need for better accuracy. Specifically, given a database containing sensitive data, we assume that a response  $y_1$  that preserves  $\epsilon_1$ -differential privacy has already been published. Then, the privacy level is relaxed to  $\epsilon_2$ , with  $\epsilon_2 > \epsilon_1$ , and we wish to publish a more accurate response  $y_2$  while the joint response  $(y_1, y_2)$  preserves  $\epsilon_2$ -differential privacy. How much accuracy is lost in the scenario of gradually releasing two responses  $y_1$  and  $y_2$  compared to the scenario of releasing a single response that is  $\epsilon_2$ -differentially private? Our results show that there exists a composite mechanism that achieves *no loss* in accuracy.

We consider the case in which the private data lies within  $\mathbb{R}^n$  with an adjacency relation induced by the  $\ell_1$ -norm, and we focus on mechanisms that approximate identity queries. We show that the same accuracy can be achieved in the case of gradual release through a mechanism whose outputs can be described by a *lazy Markov stochastic process*. This stochastic process has a closed form expression and can be efficiently sampled. Our results are applicable beyond identity queries. To this end, we demonstrate that our results can be applied in several cases, including Google's RAPPOR project, trading of sensitive data, and controlled transmission of private data in a social network. Finally, we conjecture that gradual release of data *without performance loss* is an intrinsic property of differential privacy and, thus, holds in more general settings.

---

\*Department of Electrical and Systems Engineering, University of Pennsylvania. Emails: {fkouf, hanshuo, pappasg}@seas.upenn.edu. This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

# 1 Introduction

Differential privacy is a framework that provides rigorous privacy guarantees for the release of sensitive data. The intrinsic trade-off between the privacy guarantees and accuracy of the privacy-preserving mechanism is controlled by the privacy level  $\epsilon \in [0, \infty)$ ; smaller values of  $\epsilon$  imply stronger privacy and less accuracy. Specifically, *end users*, who are interested in the output of the mechanism, demand acceptable accuracy of the privacy-preserving mechanism, whereas, *owners* of sensitive data are interested in strong enough privacy guarantees.

Existing work on differential privacy assumes that the privacy level is determined prior to release of any data and remains constant throughout the life of the privacy-preserving mechanism. However, for certain applications, the privacy level may need to be revised *after* data has been released, due to either users' need for improved accuracy or after owners' re-evaluation of the privacy concerns. One such application is trading of private data, where the owners re-evaluate their privacy concerns after monetary payments. Specifically, the end users initially access private data under  $\epsilon_1$  privacy guarantees and they later decide to "*buy*" more accurate data, relax privacy level to  $\epsilon_2$ , and enjoy better accuracy. Furthermore, the need for more accurate responses may dictate a change in the privacy level. In particular, a database containing sensitive data is persistent over time; e.g. a database of health records contains the same patients with the same health history over several years. Future uses of the database may require better accuracy, especially, after a threat is suspected (e.g. virus spread, security breach). These two example applications share the same core questions.

Is it possible to release a preliminary response with  $\epsilon_1$ -privacy guarantees and, later, release a more accurate and less private response with overall  $\epsilon_2$ -privacy guarantees? How is this scenario compared to publishing a single response under  $\epsilon_2$ -privacy guarantees? In fact, is the performance of the second response damaged by the preliminary one?

Composition theorems [1] provide a simple, but suboptimal, solution to gradually releasing sensitive data. Given an initial privacy level  $\epsilon_1$ , a noisy, privacy-preserving response  $y_1$  is generated. Later, the privacy level is increased to a new value  $\epsilon_2$  and a new response  $y_2$  is published. For an overall privacy level of  $\epsilon_2$ , the second response  $y_2$  needs to be  $(\epsilon_2 - \epsilon_1)$ -private, according to the composition theorem. Therefore, the accuracy of the second response deteriorates because of the initial release  $y_1$ .

In this work, we derive a composite mechanism which exhibits *no loss* in accuracy after the privacy level is relaxed. This mechanism employs correlation between successive responses, and, to the best of our knowledge, is the first mechanism that performs gradual release of sensitive data.

## 1.1 Our Results

This work introduces the problem of gradually releasing sensitive data. Our results focus on the case of vector-valued sensitive data  $u \in \mathbb{R}^n$  with an  $\ell_1$ -norm adjacency relation. Our first result states that, for the one-dimensional ( $n = 1$ ) identity query, there is an algorithm which relaxes privacy in two steps without sacrificing any accuracy. Although our technical treatment focuses on identical queries, our results are applicable to a broader family of queries. We also prove the *Markov property* for this algorithm and, thus, we can easily (without any computational complexity) relax privacy in any number of steps. These two results provide a different perspective of differential privacy, and lead to the definition of a *lazy* Markov stochastic process indexed by the privacy level  $\epsilon$ . Gradually releasing sensitive data is performed by sampling once from this stochastic process. We also extend the results to the high-dimensional case.

On a theoretical level, our contributions add a whole new dimension to differential privacy —

that of a *varying* parameter  $\epsilon$ . We focus on the mechanism that adds Laplace-distributed noise  $V_\epsilon$  to the private data  $u \in \mathbb{R}^n$ :

$$Q_\epsilon u = u + \begin{bmatrix} v_\epsilon^{(1)} \\ \vdots \\ v_\epsilon^{(n)} \end{bmatrix}, \quad (1)$$

where  $\epsilon$  is the privacy level,  $\|\cdot\|_1$  is the  $\ell_1$ -norm, and  $\{v_\epsilon^{(i)}\}_{i=1}^n$  are independent and identically distributed samples from the stochastic process  $\{V_\epsilon\}_{\epsilon>0}$  which has the following properties:

1.  $\{V_\epsilon\}_{\epsilon>0}$  is Markov:  $V_{\epsilon_1} \perp V_{\epsilon_3} | V_{\epsilon_2}$ , for any  $\epsilon_3 \geq \epsilon_2 \geq \epsilon_1 > 0$ .
2.  $V_\epsilon$  is Laplace-distributed:  $\mathbb{P}(V_\epsilon = x) = \frac{\epsilon}{2} e^{-\epsilon|x|}$ .
3.  $\{V_\epsilon\}_{\epsilon>0}$  is *lazy*, i.e. there is positive probability of not changing value):

$$\mathbb{P}(V_{\epsilon_1} = x | V_{\epsilon_2} = y) = \left(\frac{\epsilon_1}{\epsilon_2}\right)^2 \delta(x - y) + \left(1 - \left(\frac{\epsilon_1}{\epsilon_2}\right)^2\right) \frac{\epsilon_1}{2} e^{-\epsilon_1|x-y|}, \text{ where } \epsilon_2 \geq \epsilon_1 > 0,$$

where  $\delta$  is Dirac's delta function. For a fixed  $\epsilon$ , mechanism (1) reduces to the Laplace mechanism.

Mechanism (1) has the following properties and, thus, performs gradual release of private data:

- **Privacy:** For any set of privacy levels  $\{\epsilon_i\}_{i=1}^m$ , the mechanism that responds with  $\{Q_{\epsilon_i} u\}_{i=1}^m$  is  $(\max_{i=1}^m \epsilon_i)$ -private.
- **Accuracy:** For a fixed  $\epsilon$ , the mechanism  $Q_\epsilon$  is the optimal  $\epsilon$ -private mechanism.

In practice, gradual release of private data is achieved by sampling the stochastic process  $\{V_\epsilon\}_{\epsilon>0}$ :

1. Draw a single sample  $\{v_\epsilon\}_{\epsilon>0}$  from the stochastic process  $\{V_\epsilon\}_{\epsilon>0}$ .
2. Compute the signal  $y_\epsilon = u + v_\epsilon$ ,  $\epsilon > 0$ .
3. For  $\epsilon_1$ -privacy guarantees, release the random variable  $y_{\epsilon_1}$ .
4. Once privacy level is relaxed from  $\epsilon_1$  to  $\epsilon_2$ , where  $\epsilon_2 \geq \epsilon_1$ , release the random variable  $y_{\epsilon_2}$ .
5. In order to relax privacy level in an arbitrarily many times,  $\epsilon_1 \rightarrow \epsilon_2 \rightarrow \dots \rightarrow \epsilon_m$ , repeat the last step.

More formally, our main result derives a composite mechanism that gradually releases private data by relaxing the privacy level in an arbitrary number of steps.

**Theorem 1 (A. Gradual Privacy as a Composite Mechanism).** *Let  $\mathbb{R}^n$  be the space of privacy data equipped with an  $\ell_1$ -norm adjacency relation. Consider  $m$  privacy levels  $\{\epsilon_i\}_{i=1}^m$  such that  $0 \leq \epsilon_1 \leq \dots \leq \epsilon_m$  which successively relax the privacy level. Then, there exists a composite mechanism  $Q$  of the form*

$$Qu := (u + V_1, \dots, u + V_m), \quad (2)$$

such that:

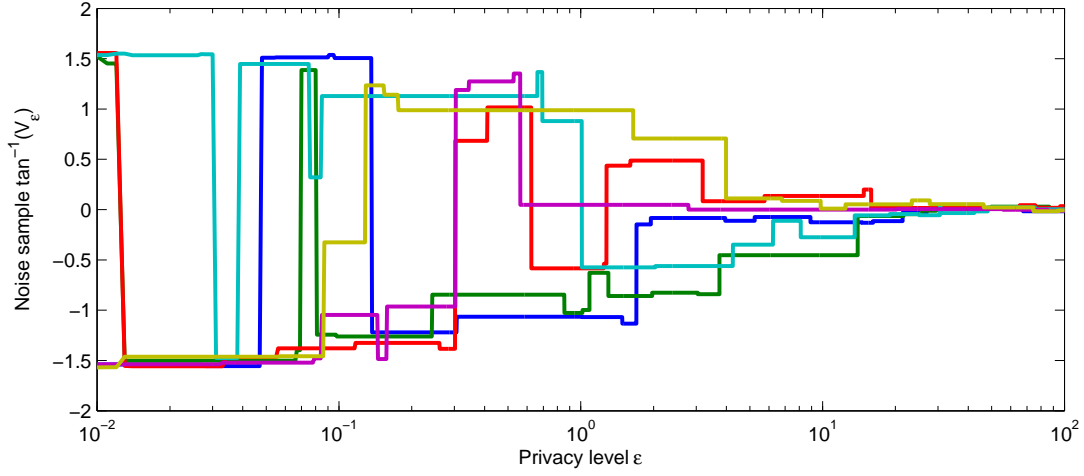


Figure 1: Gradual release of identity queries is achieved with the use of the stochastic process  $V_\epsilon$  for  $\epsilon \geq 0$ . For tight values of privacy ( $\epsilon \rightarrow 0$ ), high values of noise ( $|\tan^{-1} V_\epsilon| \rightarrow \frac{\pi}{2}$ ) are returned, whereas, almost zero samples ( $V_\epsilon \rightarrow 0$ ) are returned for large privacy budgets ( $\epsilon \rightarrow \infty$ ). The process  $V_\epsilon$  is Markov; future samples depend only on the current value of the process which eases implementation. Furthermore, the process is lazy; the value of the process changes only a few times.

1. The restriction of the mechanism  $Q$  to the first  $j$  coordinates  $(u + V_1, \dots, u + V_j)$  is  $\epsilon_j$ -private, for any  $j \in \{1, \dots, m\}$ .
2. Each coordinate  $j \in \{1, \dots, m\}$  of the mechanism  $u + V_j$  achieves the optimal mean-squared error  $\mathbb{E}\|V_j\|_2^2$ .

The mechanism that satisfies Theorem 1 has a closed-form expression and provides a new perspective of differential privacy. Instead of designing composite mechanisms of the form (2), we consider the *continuum* of privacy levels  $\epsilon \in [0, \infty)$ . Our results are more succinctly stated in terms of a stochastic process  $\{V_\epsilon\}_{\epsilon \geq 0}$ . A composite mechanism is recovered from the stochastic process by sampling the process  $\{V_\epsilon\}_{\epsilon \geq 0}$  at a finite set of privacy levels  $\{\epsilon_i\}_{i=1}^m$ .

**Theorem 1 (B. Gradual Privacy as a Stochastic Process).** *Let  $\mathbb{R}^n$  be the space of privacy data equipped with the  $\ell_1$ -norm. Then, there exists a stochastic process  $\{V_\epsilon\}_{\epsilon \geq 0}$  that defines the family of mechanisms  $Q_\epsilon$  parametrized by  $\epsilon$ :*

$$Q_\epsilon u := u + V_\epsilon, \quad \epsilon \in (0, \infty), \quad (3)$$

such that:

- **Privacy:** For any  $\epsilon > 0$ , the mechanism that releases the signal  $\{u + V_\sigma\}_{\sigma \in (0, \epsilon]}$  is  $\epsilon$ -private.
- **Accuracy:** The mechanism  $Q_\epsilon$  that releases the random variable  $u + V_\epsilon$  is the optimal  $\epsilon$ -private mechanism, i.e. the noise sample  $V_\epsilon$  achieves the optimal mean-squared error  $\mathbb{E}\|V_\epsilon\|_2^2$ .

From a more practical point of view, our results are applicable to cases beyond identity queries. Specifically, our results are directly applicable to a broad family of privacy-preserving mechanisms that are built upon the Laplace mechanism and, informally, have the following form. The sensitive data is initially preprocessed, then, the Laplace mechanism is invoked, and, finally, a post-processing

step occurs. Under the assumption that the preprocessing step is invariant of the privacy level, gradual release of sensitive data is possible. We demonstrate the applicability of our results on Google’s RAPPOR project [2], which analyzes software features that individuals use while respecting their privacy. In particular, if a software feature is suspected to be malicious, privacy level can be gradually relaxed and a more accurate analysis can be performed. On another direction, our results broaden the spectrum of applications of differential privacy. To this end, we present an application to social networks where users have different privacy concerns against close friends, acquaintances, and strangers.

We conclude our paper with a conjecture. Although present work focuses on mechanisms that add Laplace-distributed noise, we conjecture that the feasibility of gradually releasing sensitive data is a more general property of differential privacy. In particular, we formulate the conjecture that repeatedly relaxing the privacy level without loss of accuracy is possible for a larger family of privacy-aware mechanisms.

## 1.2 Previous Work

Differential privacy is an active field of research and a rich spectrum of differential private mechanisms has appeared in the literature. The exponential mechanism [1] is a powerful and generic tool for building differential private mechanisms. In particular, mechanisms that efficiently approximate linear (counting) queries have received a lot of attention [3], [4], [5]. Besides counting queries, privacy-aware versions of more complex quantities have been introduced such as signal filtering [6], optimization problems [7], [8], and allocation problems [9]. In addition to the theoretical work, differential privacy has been deployed in software tools [10].

The aforementioned work assumes that the privacy level  $\epsilon$  is a designer’s choice that is held fixed throughout the life of the privacy-aware mechanism. To the best of our knowledge, our work is the first approach that considers privacy-aware mechanisms with a varying privacy level  $\epsilon$ . Gradually releasing private data resembles the setting of differential privacy under continuous observation, which was first studied in [11]. In that setting of [11], the privacy level remains fixed while more sensitive data is being added to the database and more responses are released. In contrast, our setting assumes that both the sensitive data and the quantity of interest are fixed and the privacy level  $\epsilon$  is varying.

Gradual release of sensitive data is closely related to optimality results. Work in [3] established optimality results in an asymptotic sense (with the size of the database). Instead, our work requires *exact* optimality results and, therefore, is presented within a tighter version of differential privacy that was explored in [12], [13], where exact optimality results exist. This tighter notion which is targeted for metric spaces and we call *Lipschitz privacy*, allows for the use of optimization techniques and calculus tools. Prior work on Lipschitz privacy includes the *exact* optimality of the Laplace mechanism is established under Lipschitz privacy [13], [14].

On a more technical level, most prior work on differential privacy [7], [8], [9] introduces differential private mechanisms that are built upon the Laplace mechanism and variations of it. Although building upon the Laplace mechanism limits the solution space, there is a good reason for doing so. Specifically, for non-trivial applications, the space of probability measures can be extremely rich and hard to deal with. Technically, our approach deviates from prior work by searching over the whole space of differential private mechanisms. Work in [15] is another example that proposes a non-Laplace distribution in order to achieve better performance on subsequent queries while satisfying overall differential privacy constraints. The Laplace mechanism, then, naturally emerges as the optimal mechanism.

## 2 Background Information

### 2.1 Differential Privacy

The framework of differential privacy [16], [17] dictates that, whenever sensitive data is accessed, a noisy response is returned. The statistics of the injected noise are deliberately designed to ensure two things. First, an adversary that observes the noisy response cannot *confidently infer* the original sensitive data. The privacy level is parametrized by  $\epsilon \in [0, \infty)$ , where smaller values of  $\epsilon$  imply stronger privacy guarantees. Second, the noisy response can still be used as a surrogate of the exact response *without severe performance degradation*. On the other hand, the accuracy of the noisy response is quantified by the mean-squared error from the exact response.

Work in [16] defined differential privacy, which provides strong privacy guarantees against a powerful adversary.

**Definition 2** (Differential Privacy). *Let  $\mathcal{U}$  be a set of private data,  $\mathcal{A} \subseteq \mathcal{U}^2$  be a symmetric binary relation (called adjacency relation) and  $\mathcal{Y}$  be a set of possible responses. For  $\epsilon > 0$ , the randomized mapping  $Q : \mathcal{U} \rightarrow \Delta(\mathcal{Y})$  (called mechanism) is  $\epsilon$ -differentially private if*

$$\mathbb{P}(Qu \in \mathcal{S}) \leq e^\epsilon \mathbb{P}(Qu' \in \mathcal{S}), \forall (u, u') \in \mathcal{A}, \forall \mathcal{S} \subseteq \mathcal{Y}. \quad (4)$$

**Remark 1.** *We assume the existence of a rich-enough  $\sigma$ -algebra  $M \subseteq 2^{\mathcal{Y}}$  on the set of possible responses  $\mathcal{Y}$ . Then,  $\Delta(\mathcal{Y})$  denotes the set of probability measures over  $(M, \mathcal{Y})$ .*

Let  $y \sim Qu$  be a noisy response produced by the  $\epsilon$ -differentially private mechanism  $Q$ . For brevity, we say that “output  $y$  preserves  $\epsilon$ -privacy of the input  $u$ ”.

The adjacency relation  $\mathcal{A}$  captures the aspects of the private data  $u$  that are deemed sensitive. Consider a scheme with  $n$  users, where each user  $i$  contributes her real-valued private data  $u_i \in \mathbb{R}$ , and a private database  $u = [u_1, \dots, u_n] \in \mathbb{R}^n$  is composed. For  $\alpha > 0$ , an adjacency relation that captures the participation of a single individual to the aggregating scheme is defined as:

$$(u, u') \in \mathcal{A}_{\ell_0} \Leftrightarrow \exists j \text{ s.t. } u_i = u'_i, \forall i \neq j \text{ and } |u_j - u'_j| \leq \alpha. \quad (5)$$

Adjacency relation  $\mathcal{A}_{\ell_0}$  can be relaxed to  $\mathcal{A}_{\ell_1}$ , which is induced by the  $\ell_1$ -norm and is defined as:

$$(u, u') \in \mathcal{A}_{\ell_1} \Leftrightarrow \|u - u'\|_1 \leq \alpha, \quad (6)$$

where it holds that  $\mathcal{A}_{\ell_0} \subseteq \mathcal{A}_{\ell_1}$ .

Resilience to post-processing establishes that any post-processing on the output of an  $\epsilon$ -differentially private mechanism cannot hurt the privacy guarantees.

**Proposition 3** (Resilience to Post-Processing). *Let  $Q : \mathcal{U} \rightarrow \Delta(\mathcal{Y})$  be an  $\epsilon$ -differentially private mechanism and  $g : \mathcal{Y} \rightarrow \mathcal{Z}$  be a possibly randomized function. Then, the mechanism  $g \circ Q$  is also  $\epsilon$ -differentially private.*

More complicated mechanisms can be defined from simple ones using the composition theorem.

**Proposition 4** (Composition). *Let mechanisms  $Q_1, Q_2 : \mathcal{U} \rightarrow \Delta(\mathcal{Y})$  respectively satisfy  $\epsilon_1$  and  $\epsilon_2$ -differential privacy. Then, the composite mechanism  $Q : \mathcal{U} \rightarrow \Delta(\mathcal{Y}^2)$  defined by  $Q = (Q_1, Q_2)$  is  $(\epsilon_1 + \epsilon_2)$ -differentially private.*

Proposition 4 provides privacy guarantees whenever the *same* sensitive data is repeatedly used. Moreover, the resulting privacy level  $\epsilon_1 + \epsilon_2$  given by Proposition 4 is an upper bound and can severely over-estimate the actual privacy level. The mechanism presented in this paper introduces correlation between mechanisms  $Q_1$  and  $Q_2$ , so that it provides much stronger privacy guarantees.

## 2.2 Lipschitz Privacy

Lipschitz privacy [12], [13] is a slightly stronger version of differential privacy and is often used when the data is defined on metric spaces.

**Definition 5** (Lipschitz Privacy). *Let  $(\mathcal{U}, d)$  be a metric space and  $\mathcal{Y}$  be the set of possible responses. For  $\epsilon > 0$ , the mechanism  $Q$  is  $\epsilon$ -Lipschitz private if the following Lipschitz condition holds:*

$$|\ln \mathbb{P}(Qu \in \mathcal{S}) - \ln \mathbb{P}(Qu' \in \mathcal{S})| \leq \epsilon d(u, u'), \quad \forall u, u' \in \mathcal{U}, \forall \mathcal{S} \subseteq \mathcal{Y}. \quad (7)$$

Lipschitz privacy is closely related to the original definition of differential privacy, where the adjacency relation  $\mathcal{A}$  in differential privacy is defined through the metric  $d$ . In fact, any Lipschitz private mechanism is also differentially private.

**Proposition 6.** *For any  $\alpha > 0$ , an  $\epsilon$ -Lipschitz private mechanism  $Q$  is  $\alpha\epsilon$ -differentially private under the adjacency relation  $\mathcal{A}$ :*

$$(u, u') \in \mathcal{A} \Leftrightarrow d(u, u') \leq \alpha. \quad (8)$$

Adjacency relation  $\mathcal{A}_{\ell_1}$  defined in (6) can be captured by the  $\ell_1$ -norm under the notion of Lipschitz privacy; the metric  $d$  is  $d(u, u') = \|u - u'\|_1$ .

Our results are stated within the Lipschitz privacy framework. Proposition 6 implies that our privacy results remain valid within the framework of differential privacy. For brevity, we call an  $\epsilon$ -Lipschitz private mechanism as  $\epsilon$ -private and imply that a differentially private mechanism can be derived.

Similar to differential privacy, Lipschitz privacy is preserved under post-processing (Proposition 3) and composition of mechanisms (Proposition 4). Compared to differential privacy, Lipschitz privacy is more convenient to work with when the data and adjacency relation are defined on a metric space, which allows for the use of calculus tools. Under mild assumptions, the Lipschitz constraint (7) is equivalent to a derivative bound. In particular, for  $\mathcal{U} = \mathbb{R}^n$  equipped with the metric induced by the norm  $\|\cdot\|$ , a mechanism  $Q$  is  $\epsilon$ -Lipschitz private if

$$\|\nabla_u \ln \mathbb{P}(Qu = y)\|_* \leq \epsilon, \quad (9)$$

where  $\|\cdot\|_*$  is the dual norm of  $\|\cdot\|$ . In practice, we check condition (9) to establish the privacy properties of mechanism  $Q$ .

## 2.3 Optimality of the Laplace Mechanism

Computing the optimal private mechanism for a fixed privacy level  $\epsilon$  is considered an open problem. Specifically, let  $\mathcal{U}$  be the space of private data,  $\mathcal{A}$  be an adjacency relation,  $q : \mathcal{U} \rightarrow \mathcal{Y}$  be a query, and  $\epsilon$  be a fixed privacy level. The exponential mechanism [1] is a popular technique for constructing private mechanisms.

**Proposition 7** (Exponential Mechanism). *Let  $s : \mathcal{U} \times \mathcal{Y} \rightarrow \mathbb{R}$  be 1-Lipschitz in  $(\mathcal{U}, d)$ . Consider the mechanism  $Q$  whose output satisfies*

$$\mathbb{P}(Qu = y) \propto e^{\epsilon s(u, y)}. \quad (10)$$

*Then,  $Q$  is  $\epsilon$ -Lipschitz private.*

The Laplace mechanism is a special instance of the exponential mechanism for real spaces  $(\mathbb{R}^n, \ell_1)$ .

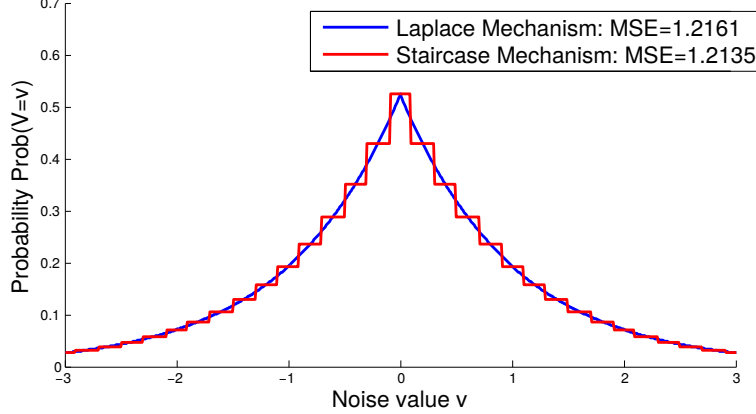


Figure 2: The staircase mechanism is the optimal  $\epsilon$ -differential private mechanism, whereas the Laplace mechanism is the optimal  $\epsilon$ -Lipschitz private mechanism. Therefore, the

**Definition 8** (Laplace Mechanism). *Let  $(\mathbb{R}^n, \ell_1)$  be the space of private data. The Laplace mechanism is defined as:*

$$Qu = u + V, \text{ where } V \sim e^{-\epsilon\|V\|_1}. \quad (11)$$

The Laplace mechanism can be shown to be  $\epsilon$ -differentially private. In general, however, the Laplace mechanism is suboptimal in the sense of minimum mean-squared error. For the single-dimensional case, the staircase mechanism [18] is the optimal  $\epsilon$ -differentially private mechanism; the mechanism which adds noise  $V$  whose distribution is shown in Figure 2. However, the Laplace mechanism is proven to be the optimal  $\epsilon$ -Lipschitz private mechanism in the sense of both minimum entropy [14] and minimum mean-squared error [13], whereas the staircase mechanism fails to satisfy Lipschitz privacy due to its discontinuous probability density function.

**Theorem 9** ([13] Optimality of Laplace). *Consider the  $\epsilon$ -Lipschitz private (in  $(\mathbb{R}^n, \ell_1)$ ) mechanism  $Q : \mathbb{R}^n \rightarrow \Delta(\mathbb{R}^n)$  of the form  $Qu = u + V$ , with  $V \sim g(V) \in \Delta(\mathbb{R}^n)$ . Then, the Laplace mechanism that adds noise with density  $l_1^n(v) = \left(\frac{\epsilon}{2}\right)^n e^{-\epsilon\|v\|_1}$  minimizes the mean-squared error. Namely, for any density  $g$ , we have:*

$$\mathbb{E}\|Qu - u\|_2^2 = \mathbb{E}_{V \sim g} \|V\|^2 \geq \mathbb{E}_{V \sim l_1^n} \|V\|_2^2 = \frac{2n}{\epsilon^2}. \quad (12)$$

The optimal private mechanism characterizes the privacy-performance trade-off and is required for gradually releasing sensitive data. Thus, optimality of the Laplace mechanism in Theorem 9 is a key ingredient in our results and renders the problem tractable.

### 3 Gradual Release of Private Data

The problem of gradually releasing private data is now formulated. Initially, we focus on a single privacy level relaxation from  $\epsilon_1$  to  $\epsilon_2$  and a single-dimensional space of private data  $\mathcal{U} = \mathbb{R}$ . Subsections 3.2 and 3.3 present extensions to high-dimensional spaces and multiple rounds of privacy level relaxations, respectively.

Consider two privacy levels  $\epsilon_1$  and  $\epsilon_2$  with  $\epsilon_2 \geq \epsilon_1 > 0$ . We wish to design a composite mechanism  $Q_{\epsilon_1 \rightarrow \epsilon_2} : \mathcal{U} \rightarrow \Delta(\mathcal{Y} \times \mathcal{Y})$  that performs gradual release of data. The first and second



coordinates respectively refer to the initial  $\epsilon_1$ -private and the subsequent  $\epsilon_2$ -private responses. In practice, given privacy levels  $\epsilon_1$  and  $\epsilon_2$  and an input  $u \in \mathcal{U}$ , we sample  $(y_1, y_2)$  from the distribution  $Q_{\epsilon_1 \rightarrow \epsilon_2} u$ . Initially, only coordinate  $y_1$  is published satisfying  $\epsilon_1$ -privacy guarantees. Once privacy level is relaxed to  $\epsilon_2$ , response  $y_2$  is released as a more accurate response of the *same* query on the *same* private data.

An adversary that wishes to infer the private input  $u$  eventually has access to both responses  $y_1$  and  $y_2$ . Therefore, the pair  $(y_1, y_2)$  needs to satisfy  $\epsilon_2$ -privacy. On the other hand, an honest user wishes to maximize the accuracy of the response and, therefore, she is tempted to use an estimator  $y_M = \theta(y_1, y_2)$  and infer a more accurate response  $y_M$ . In order to relieve honest users from any computational burden, we wish the best estimator to be as the truncation:

$$\theta(y_1, y_2) = y_2. \quad (13)$$

The composition theorem [1] provides a trivial, yet highly conservative, approach. Specifically, compositional rules imply that, if  $y_1$  satisfies  $\epsilon_1$ -privacy and  $(y_1, y_2)$  satisfies  $\epsilon_2$ -privacy, coordinate  $y_2$  itself should be  $(\epsilon_2 - \epsilon_1)$ -private. In the extreme case that  $\epsilon_2 - \epsilon_1 = \delta \ll 1$ , response  $y_2$  alone is expected to be  $\delta$ -private and, therefore, is highly corrupted by noise. This is unacceptable, since estimator (13) yields an even noisier response than the initial response  $y_1$ . Even if honest users are expected to compute more complex estimators than the truncation one in (13), the approach dictated by composition theorem can still be unsatisfactory.

Specifically, consider the following two scenarios:

1. An  $\epsilon_1$ -private response  $y_1$  is initially released. Once privacy level is relaxed from  $\epsilon_1$  to  $\epsilon_2$ , an supplementary response  $y_2$  is released.
2. No response is initially released. Response  $\hat{y}_2$  is released as soon as the privacy level is relaxed to  $\epsilon_2$ .

Then, there is no guarantee that the best estimator  $\theta(y_1, y_2)$  in Scenario 1 will match the accuracy of the response  $\hat{y}_2$  in Scenario 2. An accuracy gap between the two scenarios would severely impact differential privacy. Specifically, the system designer needs to be strategic when choosing a privacy level. Differently stated, a market of private data based on composition theorems would exhibit *friction*.

The key idea to overcome this friction is to introduce correlation between responses  $y_1$  and  $y_2$ . In this work, we focus on Euclidean spaces  $\mathcal{U} = \mathbb{R}^n$  and mechanisms  $Qu = u + V$  that approximate the identity query  $q(u) = u$ . Our main result states that a *frictionless* market of private data is feasible and Scenarios 1 and 2 are equivalent. This result has multi-fold implications:

- A system designer is not required to be strategic with the choice of the privacy level. Specifically, she can initially under-estimate the required privacy level with  $\epsilon_1$  and she can later fine-tune it to  $\epsilon_2$  without hurting the accuracy of the final response.
- A privacy data market can exist and private data can be traded “*by the pound*”. An  $\epsilon_1$ -private response  $y_1$  can be initially purchased. Next, a supplementary payment can be made in return for a privacy level relaxation to  $\epsilon_2$  and a refined response  $y_2$ . The accuracy of the refined response  $y_2$  is, then, unaffected by the initial transaction and is controlled only by the final privacy level  $\epsilon_2$ .

More concretely, given privacy levels  $\epsilon_1$  and  $\epsilon_2$  with  $\epsilon_2 > \epsilon_1$ , we wish to design a composite mechanism  $Q_{\epsilon_1 \rightarrow \epsilon_2} : \mathcal{U} \rightarrow \mathcal{Y} \times \mathcal{Y}$  with the following properties:

1. The restriction of  $Q_{\epsilon_1 \rightarrow \epsilon_2}$  to the first coordinate should match the performance of the optimal  $\epsilon_1$ -private mechanism  $Q_{\epsilon_1}$ . More restrictively, the first coordinate of the composite mechanism  $Q_{\epsilon_1 \rightarrow \epsilon_2}$  should be distributed identically to the optimal  $\epsilon_1$ -private mechanism  $Q_{\epsilon_1}$ :

$$\mathbb{P}(Q_{\epsilon_1 \rightarrow \epsilon_2} u \in \mathcal{S} \times \mathcal{Y}) = \mathbb{P}(Q_{\epsilon_1} u \in \mathcal{S}), \forall u \in \mathcal{U} \text{ and } \mathcal{S} \subseteq \mathcal{Y} \quad (14)$$

2. The restriction of  $Q_{\epsilon_1 \rightarrow \epsilon_2}$  to the first coordinate should be  $\epsilon_1$ -private. This property is imposed by constraint 1.
3. The restriction of  $Q_{\epsilon_1 \rightarrow \epsilon_2}$  to the second coordinate should match the performance of the optimal  $\epsilon_2$ -private mechanism  $Q_{\epsilon_2}$ . Similarly to the first coordinate, the second coordinate of the composite mechanism  $Q_{\epsilon_1 \rightarrow \epsilon_2}$  must be distributed identically to the optimal  $\epsilon_2$ -private mechanism  $Q_{\epsilon_2}$ :

$$\mathbb{P}(Q_{\epsilon_1 \rightarrow \epsilon_2} u \in \mathcal{Y} \times \mathcal{S}) = \mathbb{P}(Q_{\epsilon_2} u \in \mathcal{S}), \forall u \in \mathcal{U} \text{ and } \mathcal{S} \subseteq \mathcal{Y} \quad (15)$$

4. Once both coordinates are published,  $\epsilon_2$ -privacy should be guaranteed. According to Lipschitz privacy, the requirement is stated as follows:

$$\mathbb{P}(Q_{\epsilon_1 \rightarrow \epsilon_2} u \in \mathcal{S}) \text{ is } \epsilon_2\text{-Lipschitz in } u, \text{ for all } \mathcal{S} \subseteq \mathcal{Y}^2. \quad (16)$$

Equations (14) and (15) require knowledge of the optimal  $\epsilon$ -private mechanism. In general, computing the  $\epsilon$ -private mechanism that maximizes a reasonable performance criterion is still an open problem. Theorem 9 establish the optimality of the Laplace mechanism as the optimal private approximation of the identity query.

### 3.1 Single-Dimensional Case

Initially, we consider the single-dimensional case where  $\mathcal{U} = \mathbb{R}$  equipped with the absolute value. Theorem 9 establish the optimal  $\epsilon$ -private mechanism that is required by Equations (14) and (15):

$$Q_\epsilon u = u + V, \text{ where } V \sim e^{-\epsilon|V|}. \quad (17)$$

Mechanism (17) minimizes the mean-squared error from the identity query among all  $\epsilon$ -private mechanisms that use additive noise:

$$\mathbb{E}_{V \sim e^{-\epsilon|V|}} (Q_\epsilon u - u)^2 \quad (18)$$

Theorem 10 establishes the existence of a composite mechanism that relaxes privacy from  $\epsilon_1$  to  $\epsilon_2$  without any loss of performance.

**Theorem 10.** Consider privacy levels  $\epsilon_1$  and  $\epsilon_2$  with  $\epsilon_2 \geq \epsilon_1 > 0$ , and mechanisms of the form:

$$Q_1 u := u + V_1 \text{ and } Q_2 u := u + V_2, \text{ with } (V_1, V_2) \sim g \in \Delta(\mathbb{R}^2). \quad (19)$$

Then, for density  $l_{\epsilon_1, \epsilon_2}$  with:

$$l_{\epsilon_1, \epsilon_2}(x, y) = \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} \delta(x - y) + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} e^{-\epsilon_1|x-y| - \epsilon_2|y|}, \quad (20)$$

where  $\delta$  is the Dirac delta function, the following properties hold:

1. The mechanism  $Q_1$  is  $\epsilon_1$ -private.
2. The mechanism  $Q_1$  is optimal, i.e.  $Q_1$  minimizes the mean-squared error  $\mathbb{E}V_1^2$ .
3. The mechanism  $(Q_1, Q_2)$  is  $\epsilon_2$ -private.
4. The mechanism  $Q_2$  is optimal, i.e.  $Q_2$  minimizes the mean-squared error  $\mathbb{E}V_2^2$ .

*Proof.* Consider the mechanism  $Q = (Q_1, Q_2)$  induced by the noise density (20). We prove that this mechanism satisfies all the desired properties:

1. The first coordinate is Laplace-distributed with parameter  $\frac{1}{\epsilon_1}$ . For  $x \geq 0$ , we get:

$$\begin{aligned}
\mathbb{P}(V_1 = x) &= \int_{\mathbb{R}} g(x, y) dy = \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2 x} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \int_{\mathbb{R}} e^{-\epsilon_1|x-y| - \epsilon_2|y|} dy \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2 x} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \left( \int_{-\infty}^0 e^{-\epsilon_1 x + (\epsilon_1 + \epsilon_2)y} dy + \int_0^x e^{-\epsilon_1 x - (\epsilon_2 - \epsilon_1)y} dy \right. \\
&\quad \left. + \int_x^{\infty} e^{\epsilon_1 x - (\epsilon_1 + \epsilon_2)y} dy \right) \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2 x} + \frac{\epsilon_1(\epsilon_2 - \epsilon_1)}{4\epsilon_2} e^{-\epsilon_1 x} e^{(\epsilon_1 + \epsilon_2)y} \Big|_{-\infty}^0 - \frac{\epsilon_1(\epsilon_1 + \epsilon_2)}{4\epsilon_2} e^{-\epsilon_1 x} e^{-(\epsilon_2 - \epsilon_1)y} \Big|_0^x \\
&\quad - \frac{\epsilon_1(\epsilon_2 - \epsilon_1)}{4\epsilon_2} e^{-(\epsilon_1 + \epsilon_2)y} \Big|_x^{\infty} \\
&= \frac{\epsilon_1}{2} e^{-\epsilon_1 x}
\end{aligned} \tag{21}$$

The case  $x \leq 0$  follows from the symmetry  $(x, y) \rightarrow (-x, -y)$ . Therefore, the first coordinate is  $\epsilon_1$ -private and achieves optimal performance.

2. The second coordinate is Laplace-distributed with parameter  $\frac{1}{\epsilon_2}$ . We have:

$$\begin{aligned}
\mathbb{P}(V_2 = y) &= \int_{\mathbb{R}} g(x, y) dx = \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} e^{-\epsilon_2|y|} \int_{\mathbb{R}} e^{-\epsilon_1|x-y|} dx \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} e^{-\epsilon_2|y|} \int_{\mathbb{R}} e^{-\epsilon_1|x|} dx \\
&= \frac{\epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} + \frac{\epsilon_2^2 - \epsilon_1^2}{2\epsilon_2} e^{-\epsilon_2|y|} \\
&= \frac{\epsilon_2}{2} e^{-\epsilon_2|y|}
\end{aligned} \tag{22}$$

Thus, the second coordinate achieves optimal performance.

3. Lastly, we need to prove that the composite mechanism is  $\epsilon_2$ -private. We handle the delta part separately by defining  $D = \{x : (x, x) \in S\}$  for a measurable  $S \subseteq \mathbb{R}^2$ . The probability of landing in set  $S$  is:

$$\mathbb{P}(Qu \in S) = \frac{\epsilon_1^2}{2\epsilon_2} \int_D e^{-\epsilon_2|x-u|} dx + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \iint_S e^{-\epsilon_1|(x-u)-(y-u)| - \epsilon_2|y-u|} dx dy \tag{23}$$

We take the derivative and use Fubini's theorem to exchange the derivative with the integral:

$$\begin{aligned}
\frac{d}{du}\mathbb{P}(Qu \in S) &= \frac{\epsilon_1^2}{2\epsilon_2} \int_D \epsilon_2 \text{sgn}(x-u) e^{-\epsilon_2|x-u|} dx \\
&\quad + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \iint_S \epsilon_2 \text{sgn}(y-u) e^{-\epsilon_1|x-y|-\epsilon_2|y-u|} dx dy \Rightarrow \\
\left| \frac{d}{du}\mathbb{P}(Qu \in S) \right| &\leq \frac{\epsilon_1^2}{2\epsilon_2} \int_D \epsilon_2 e^{-\epsilon_2|x-u|} dx + \frac{\epsilon_1(\epsilon_2^2 - \epsilon_1^2)}{4\epsilon_2} \iint_S \epsilon_2 e^{-\epsilon_1|(x-u)-(y-u)|-\epsilon_2|y-u|} dx dy \Rightarrow \\
\left| \frac{d}{du}\mathbb{P}(Qu \in S) \right| &\leq \epsilon_2 \mathbb{P}(Qu \in S) \Rightarrow \left| \frac{d}{du} \ln \mathbb{P}(Qu \in S) \right| \leq \epsilon_2
\end{aligned} \tag{24}$$

This completes the proof.  $\square$

### 3.1.1 Single Round of Privacy Relaxation

Theorem 10 achieves gradual release of sensitive data in two steps, first with  $\epsilon_1$ -privacy and, then, with  $\epsilon_2$ -privacy. In practice, Theorem 10 can be used as follows:

- Given the private value  $u \in \mathbb{R}$ , sample noise  $V_1 \sim e^{-\epsilon_1|V_1|}$  and release response  $y_1 = u + V_1$ , which is optimal and respects  $\epsilon_1$ -privacy.
- Once privacy level is relaxed from  $\epsilon_1$  to  $\epsilon_2$ , sample noise  $V_2$  from the conditioned on  $V_1$  distribution:

$$\mathbb{P}(V_2 = y | V_1 = x) = \frac{\epsilon_1}{\epsilon_2} e^{-(\epsilon_2 - \epsilon_1)|x|} \delta(y - x) + \frac{\epsilon_2^2 - \epsilon_1^2}{2\epsilon_2} e^{-\epsilon_1|y-x|-\epsilon_2|y|+\epsilon_1|x|}, \tag{25}$$

and release response  $y_2 = u + V_2$ . Distribution (25) is derived from the joint distribution (20) and ensures both that  $(y_1, y_2)$  is  $\epsilon_2$ -private and that  $V_2$  is optimally distributed.

Conditional distribution (25) is shown in Figure 3. Note that for  $\epsilon_2 = \epsilon_1$ , Distribution (25) is reduced to a delta function:

$$\mathbb{P}(V_2 = y | V_1 = x) = \delta(x - y). \tag{26}$$

In words, for  $\epsilon_2 = \epsilon_1$  no privacy relaxation effectively happens and, thus, no updated response is practically released. Moreover, for  $\epsilon_2 \rightarrow \infty$ , a limiting argument shows that Distribution 25 is reduced to:

$$\mathbb{P}(V_2 = y | V_1 = x) = \delta(y). \tag{27}$$

Practically, letting  $\epsilon_2 \rightarrow \infty$  cancel any privacy constraints and the exact value of private data  $u$  can be released  $y_2 = u$ . For general values of  $\epsilon_1$  and  $\epsilon_2$ , Pearson's correlation coefficient decreases for more aggressive privacy relations,  $\rho_{V_1, V_2} = \frac{\epsilon_1}{\epsilon_2}$ . Algorithm 1 provides a simple and efficient way to sample  $V_2$  given  $V_1$ .




Figure 3: Gradual release of private data is performed in the following way. First, the  $\epsilon_1$ -private response  $y_1 = u + V_1$  is released, where  $V_1 \sim e^{-\epsilon_1|V_1|}$ . Once privacy level is relaxed from  $\epsilon_1 = 1$  to  $\epsilon_2 = 2$ , the supplementary response  $y_2 = u + V_2$  is released, where  $V_2$  is distributed as shown above. The composite mechanism that releases  $(y_1, y_2)$  is  $\epsilon_2$ -private and  $V_2$  is optimally distributed.

---

**Algorithm 1** Sampling from Distribution (25) for the second noise sample  $V_2 = y$  given the first noise sample  $V_1 = x$  can be efficiently performed.

---

**Require:** Privacy levels  $\epsilon_1$  and  $\epsilon_2$ , such that  $\epsilon_2 > \epsilon_1 > 0$ , and noise sample  $x$ .

```

function RELAXPRIVACY( $x, \epsilon_1, \epsilon_2$ )
  switch randomly
    case with probability  $\frac{\epsilon_1}{\epsilon_2} e^{-(\epsilon_2 - \epsilon_1)|x|}$ :
      return  $y = x$ .
    case with probability  $\frac{\epsilon_2 - \epsilon_1}{2\epsilon_2}$ :
      draw  $z \sim \begin{cases} e^{(\epsilon_1 + \epsilon_2)z}, & \text{for } z \leq 0 \\ 0, & \text{otherwise.} \end{cases}$ 
      return  $y = \text{sgn}(x)z$ .
    case with probability  $\frac{\epsilon_1 + \epsilon_2}{2\epsilon_2} (1 - e^{-(\epsilon_2 - \epsilon_1)|x|})$ :
      draw  $z \sim \begin{cases} e^{-(\epsilon_2 - \epsilon_1)z}, & \text{for } 0 \leq z \leq |x| \\ 0, & \text{otherwise.} \end{cases}$ 
       $y = \text{sgn}(x)z$ .
    case with probability  $\frac{\epsilon_2 - \epsilon_1}{2\epsilon_2} e^{-(\epsilon_2 - \epsilon_1)|x|}$ :
      draw  $z \sim \begin{cases} e^{-(\epsilon_1 + \epsilon_2)z}, & \text{for } z \geq |x| \\ 0, & \text{otherwise.} \end{cases}$ 
      return  $y = \text{sgn}(x)z$ .
  end switch
end function

```

---

### 3.1.2 Single Round of Privacy Tightening

Tightening the privacy level is impossible, since it implies revoking already released data. Nonetheless, generating a more private version of the same data is still useful in cases such as *private data trading*. In that case, distribution (20) can be sampled in the opposite direction. Specifically, noise  $V_2$  is initially sampled,  $V_2 \sim e^{-\epsilon_2|V_2|}$ , and the  $\epsilon_2$ -private response  $y_2 = u + V_2$  is released. Next, private data  $u$  is traded to a *different agent* under the stronger  $\epsilon_1$ -privacy guarantees. Noise sample  $V_1$  is drawn from distribution

$$\mathbb{P}(V_1 = x | V_2 = y) = \left(\frac{\epsilon_1}{\epsilon_2}\right)^2 \delta(x - y) + \left(1 - \left(\frac{\epsilon_1}{\epsilon_2}\right)^2\right) \frac{\epsilon_1}{2} e^{-\epsilon_1|x-y|}, \quad (28)$$

and the  $\epsilon_1$ -private response  $y_1 = u + V_1$  is released. Remarkably, response  $y_1$  can be generated conditioning only on  $y_2$ :

$$y_2 = y_1 + V_{2 \rightarrow 1}, \quad (29)$$

where  $V_{2 \rightarrow 1} = V_1 - V_2$  is independent of  $V_2$ ,  $V_{2 \rightarrow 1} \perp V_2$ . In words, tightening privacy under the Laplace mechanism does not require access to the original data  $u$  and can be performed by an agent other than the private data owner. Theorem 3 suggests that the randomized post-processing  $y_1 = y_2 + V_{2 \rightarrow 1}$  of the  $\epsilon_2$ -private response  $y_2$  is at least  $\epsilon_2$ -private. For  $V_{2 \rightarrow 1}$  given by distribution (28), this tightening of privacy level is precisely quantified, i.e.  $\epsilon_2 \rightarrow \epsilon_1$ . Recall that our results are *tight*; no excessive accuracy is sacrificed in the process.

## 3.2 High-Dimensional Case

Theorem 10 can be generalized for the case that the space of private data is Euclidean  $\mathbb{R}^n$  equipped with the  $\ell_1$ -norm. Theorem 9 establishes that the Laplace mechanism:

$$Q_\epsilon u = u + V, \text{ where } V \sim e^{-\epsilon\|V\|_1}. \quad (30)$$

minimizes the mean-squared error from the identity query among all  $\epsilon$ -private mechanisms that use additive noise  $V \in \mathbb{R}^n$ :

$$\mathbb{E}_{V \sim e^{-\epsilon\|V\|_1}} \|Q_\epsilon u - u\|_2^2. \quad (31)$$

Theorem 9 shows that each coordinate of  $V$  is independently sampled. This observation implies that Theorem 10 can be applied to  $n$  dimensions independently.

**Theorem 11.** *Consider privacy levels  $\epsilon_1, \epsilon_2$  with  $\epsilon_2 > \epsilon_1 > 0$ . Let  $Q_1$  be an  $\epsilon_1$ -private mechanism and  $Q_2$  an  $\epsilon_2$ -private mechanism of the form:*

$$Q_1 u := u + V_1 \text{ and } Q_2 u := u + V_2 \text{ with } (V_1, V_2) \sim g \in \Delta(\mathbb{R}^{2n}), \quad (32)$$

where  $u \in \mathbb{R}_{\ell_1}^n$ . Then, gradual release of sensitive data  $u$  from  $\epsilon_1$  to  $\epsilon_2$  is achieved by the probability distribution  $l_{\epsilon_1, \epsilon_2}^n$ :

$$l_{\epsilon_1, \epsilon_2}^n(V_1, V_2) = \prod_{i=1}^n l_{\epsilon_1, \epsilon_2}(V_1^{(i)}, V_2^{(i)}), \quad (33)$$

where  $V_i = [V_i^{(1)}, \dots, V_i^{(n)}]$ ,  $i = 1, 2$ . Namely:

- Mechanism  $Q_1$  is  $\epsilon_1$ -private and optimal.
- Mechanism  $Q_2$  is the optimal  $\epsilon_2$ -private mechanism.
- Mechanism  $(Q_1, Q_2)$  is  $\epsilon_2$ -private.

*Proof.* Let  $[x^{(1)}, \dots, x^{(n)}]$  denote the coordinates of a vector  $x \in \mathbb{R}^n$ . The desired probability distribution is defined by independently sampling each coordinate using Theorem 10. Let:

$$l_{\epsilon_1, \epsilon_2}^n(x, y) = g(x, y) = \prod_{i=1}^n l_{\epsilon_1, \epsilon_2}(x^{(i)}, y^{(i)}), \quad (34)$$

The probability distribution satisfies the required marginal distributions:

$$\int_{\mathbb{R}^n} g(x, y) d^n y = \left(\frac{\epsilon_1}{2}\right)^n e^{-\epsilon_1 \|x\|_1} \quad \text{and} \quad \int_{\mathbb{R}^n} g(x, y) d^n x = \left(\frac{\epsilon_2}{2}\right)^n e^{-\epsilon_2 \|y\|_1}$$

Moreover, it satisfies  $\epsilon_2$ -privacy constraints:

$$\begin{aligned} & \|\nabla_u \ln \mathbb{P}(Q_1 u = z_1 \text{ and } Q_2 u = z_2)\|_\infty \\ &= \|\nabla_u \ln l_{\epsilon_1, \epsilon_2}^n(z_1 - u, z_2 - u)\|_\infty \\ &= \max_{1 \leq i \leq n} \left| \frac{\partial}{\partial u_i} \ln l_{\epsilon_1, \epsilon_2}^n(z_1 - u, z_2 - u) \right| \\ &= \max_{1 \leq i \leq n} \left| \frac{\partial}{\partial u_i} \ln l_{\epsilon_1, \epsilon_2}(z_1^{(i)} - u^{(i)}, z_2^{(i)} - u^{(i)}) \right| \\ &\leq \max_{1 \leq i \leq n} \epsilon_2 = \epsilon_2, \end{aligned}$$

where in the last line we used the fact that  $l_{\epsilon_1, \epsilon_2}$  is  $\epsilon_2$ -private. This completes the proof.  $\square$

### 3.3 Multiple Privacy Relaxations

Theorems 10 and 11 perform privacy relaxation from  $\epsilon_1$  to  $\epsilon_2$ . However, the privacy level is possibly updated multiple times. Theorem 12 handles the case where the privacy level is successively relaxed from  $\epsilon_1$  to  $\epsilon_2$ , to  $\epsilon_3$ , until  $\epsilon_m$ . Specifically, Theorem 12 enables the use of Theorem 10 multiple times while relaxing privacy level from  $\epsilon_i$  to  $\epsilon_{i+1}$  for  $i \in \{1, \dots, m-1\}$ . We call this statement the *Markov property* of the Laplace mechanism.

**Theorem 12.** Consider  $m$  privacy levels  $\{\epsilon_i\}_{i=1}^m$  with  $0 < \epsilon_1 < \dots < \epsilon_m$  and mechanisms  $Q_i$  of the form:

$$Q_i u = u + V_i, \text{ with } (V_1, \dots, V_m) \sim g \in \Delta(\mathbb{R}^m). \quad (35)$$

Consider the distribution  $g = l_{\epsilon_1, \dots, \epsilon_m}$ , with:

$$l_{\epsilon_1, \dots, \epsilon_m}(v_1, \dots, v_m) = l_{\epsilon_1}(v_1) \prod_{i=1}^{m-1} \frac{l_{\epsilon_i, \epsilon_{i+1}}(v_i, v_{i+1})}{l_{\epsilon_i}(v_i)}, \quad (36)$$

where  $l_\epsilon(v) = \frac{\epsilon}{2} e^{-\epsilon|v|}$ . Then, distribution  $l_{\epsilon_1, \dots, \epsilon_m}$  has the following properties:

1. Each prefix mechanism  $(Q_1, \dots, Q_i)$  is  $\epsilon_i$ -private, for  $i \in \{1, \dots, m\}$ .

2. Each mechanism  $Q_i$  is the optimal  $\epsilon_i$ -private mechanism, i.e. it minimizes the mean-squared error  $\mathbb{E}V_i^2$ .

*Proof.* The proof uses induction on  $m$ . The case  $m = 2$  is handled by Theorem 10. For brevity, we prove the statement for  $m = 3$ . Let  $f(x, y) = l_{\epsilon_1, \epsilon_2}(x, y)$  and  $g(y, z) = l_{\epsilon_2, \epsilon_3}(y, z)$ . Consider the joint probability  $l_{\epsilon_1, \epsilon_2, \epsilon_3}$ :

$$h(x, y, z) = l_{\epsilon_1, \epsilon_2, \epsilon_3}(x, y, z) = \frac{f(x, y)g(y, z)}{l_{\epsilon_2}(y)}, \quad (37)$$

where  $l_{\epsilon_2}(y) = \frac{\epsilon_2}{2} e^{-\epsilon_2|y|}$ . Measure  $h$  possesses all the properties that perform gradual release of private data:

- All marginal distributions of measure  $h$  are Laplace with parameters  $\frac{1}{\epsilon_1}$ ,  $\frac{1}{\epsilon_2}$ , and  $\frac{1}{\epsilon_3}$ , respectively:

$$\int_{\mathbb{R}} \int_{\mathbb{R}} h(x, y, z) dy dz = l_{\epsilon_1}(x), \quad \int_{\mathbb{R}} \int_{\mathbb{R}} h(x, y, z) dx dz = l_{\epsilon_2}(y), \quad \text{and} \quad \int_{\mathbb{R}} \int_{\mathbb{R}} h(x, y, z) dx dy = l_{\epsilon_3}(z).$$

- Mechanism  $Q_1$  is  $\epsilon_1$ -private since  $V_1$  is Laplace-distributed with parameter  $\frac{1}{\epsilon_1}$ .
- Mechanism  $(Q_1, Q_2)$  is  $\epsilon_2$ -private. Margining out  $V_3$  shows that  $(V_1, V_2) \sim l_{\epsilon_1, \epsilon_2}$ , which guarantees  $\epsilon_2$ -privacy according to Theorem 10.
- Mechanism  $(Q_1, Q_2, Q_3)$  is  $\epsilon_3$ -private. It holds that:

$$\begin{aligned} & \left| \frac{\partial}{\partial u} \mathbb{P}(Q_1 u = \psi_1, Q_2 u = \psi_2, \text{ and } Q_3 u = \psi_3) \right| \\ &= \left| \frac{\partial}{\partial u} h(\psi_1 - u, \psi_2 - u, \psi_3 - u) \right| \\ &= \left| \frac{\partial h(x, y, z)}{\partial x} + \frac{\partial h(x, y, z)}{\partial y} + \frac{\partial h(x, y, z)}{\partial z} \right| \Bigg|_{\substack{x=\psi_1-u, \\ y=\psi_2-u, \\ z=\psi_3-u}} \end{aligned} \quad (38)$$

Algebraic manipulation of the last expression establishes the result:

$$\begin{aligned} \left| \frac{\partial h}{\partial x} + \frac{\partial h}{\partial y} + \frac{\partial h}{\partial z} \right| &= \left| \frac{f_x g}{l_{\epsilon_2}} + \frac{f_y g}{l_{\epsilon_2}} + \frac{f g_y}{l_{\epsilon_2}} - l_{\epsilon_2} \frac{f g}{l_{\epsilon_2}^2} + \frac{f g_z}{l_{\epsilon_2}} \right| \\ &= \left| -\text{sgn}(y) \epsilon_2 \frac{f g}{l_{\epsilon_2}} - \text{sgn}(z) \epsilon_3 \frac{f g}{l_{\epsilon_2}} + \text{sgn}(y) \epsilon_2 \frac{f g}{l_{\epsilon_2}} \right| \\ &= \left| -\text{sgn}(z) \epsilon_3 \frac{f g}{l_{\epsilon_2}} \right| \\ &= \epsilon_3 h, \end{aligned} \quad (39)$$

where we used the properties  $l'_{\epsilon_2} = -\text{sgn}(y) \epsilon_2 l_{\epsilon_2}$ ,  $f_x + f_y = -\text{sgn}(y) \epsilon_2 f$ , and  $g_y + g_z = -\text{sgn}(z) \epsilon_3 g$ , where the last two identities were derived in the proof of Theorem 10.  $\square$

Additionally, performing multiple rounds of privacy relaxations can be performed in the context of Theorem 11 is possible. In that case, Theorem 12 is independently applied to each component.



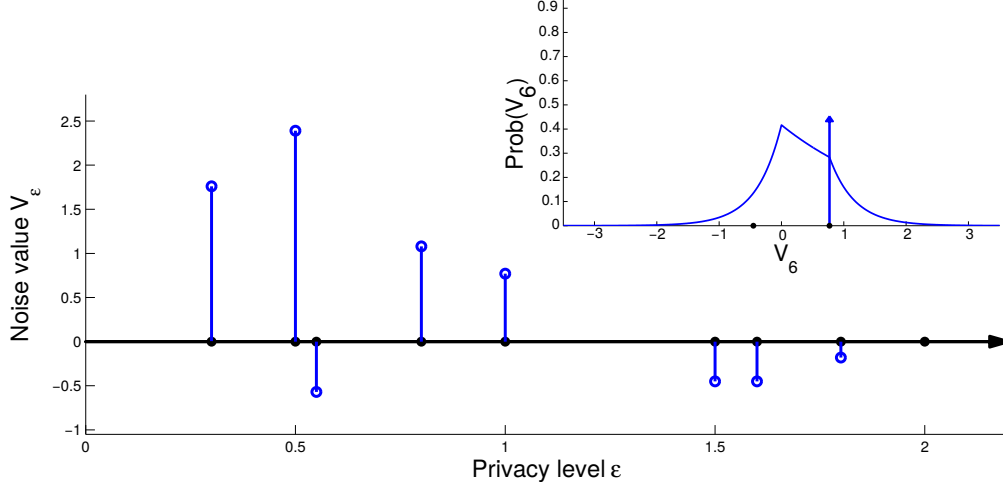


Figure 4: Privacy level can be repeatedly relaxed. For each round of relaxation  $\epsilon_i \rightarrow \epsilon_{i+1}$ , the distribution of the next noise sample  $V_{i+1}$  depends only on the last noise sample  $V_i$ . Past noise samples  $\{V_j\}_{j < i}$  can be discarded from memory, thus, there is no complexity incurred from repeatedly relaxing privacy level.

### 3.3.1 Multiple Rounds of Privacy Relaxation

Theorem 12 states that it is possible to repeatedly use Theorem 10 to perform multiple privacy level relaxations. An intuitive proof of Theorem 12 can be constructed by considering Scenarios 1 and 2 introduced in the beginning of Section 3. Specifically, Theorem 10 constructs a *coupling* such that Scenario 1 replicates Scenario 2. Therefore, once the first round of privacy relaxation  $\epsilon_1 \rightarrow \epsilon_2$  occurs, the two scenarios are indistinguishable. The second round of privacy relaxation  $\epsilon_2 \rightarrow \epsilon_3$  is performed by starting at the first step of Scenario 1.

In practice, Theorem 12 allows for an efficient implementation of an arbitrary number of privacy relaxation rounds  $\epsilon_1 \rightarrow \epsilon_2 \rightarrow \dots \rightarrow \epsilon_m$ . In particular, only the most recent privacy level  $\epsilon_i$  and noise sample  $V_i$  need to be stored in memory. Sampling for  $V_{i+1}$  depends only on current privacy level  $\epsilon_i$ , current noise sample  $V_i$  and next privacy level  $\epsilon_{i+1}$ . Past privacy levels  $\{\epsilon_j\}_{j < i}$ , past noise samples  $\{V_j\}_{j < i}$ , and future privacy levels  $\{\epsilon_j\}_{j > i+1}$  are not needed.

## 3.4 A Private Stochastic Process

Theorems 10 and 12 offer a novel dimension to the Laplace mechanism. Specifically, these results establish a real-valued stochastic process  $\{V_\epsilon : \epsilon > 0\}$ . Sampling from the process  $\{V_\epsilon\}_{\epsilon > 0}$  performs gradual release of sensitive data for the continuum of privacy levels  $(0, \infty)$ . Consider the mechanisms  $Q_\epsilon$  that respond with  $Q_\epsilon u = y_\epsilon = u + V_\epsilon$ . Then:

- $V_\epsilon$  is optimally distributed, i.e. Laplace-distributed with parameter  $\frac{1}{\epsilon}$
- Any  $\epsilon$ -truncated response  $\{y_\sigma\}_{\sigma \in (0, \epsilon]}$  is  $\epsilon$ -private.

Samples of the process  $\{V_\epsilon\}_{\epsilon > 0}$  are plotted in Figure 1. This process features properties that allow efficient sampling:

- It is *Markov*,  $V_s \perp V_t | V_q$ , for  $s < q < t$ . Thus, a sample of the process  $V_\epsilon$  over an interval  $[\epsilon_1, \epsilon_2]$  can be extended to  $[\epsilon_1, \epsilon_3]$ , for  $\epsilon_3 > \epsilon_2$ .

- It is *lazy*, i.e.  $V_{\epsilon+\delta} = V_\epsilon$  with high probability, for  $\delta \ll 1$ . Therefore, A sample of the process  $\{V_\epsilon\}_{\epsilon_1 \leq \epsilon \leq \epsilon_2}$  can be efficiently stored; only a finite (random) number  $m$  of points  $(\epsilon_i, V_{\epsilon_i})_{i=1}^m$  where jumps occur need to be stored for *exact* re-construction of the process.

## 4 Applications

### 4.1 Crowdsourcing Statistics with RAPPOR

Theorems 10 and 12 perform gradual release of private data by releasing responses that approximate the identity query  $q(u) = u$ . In practice, however, the end-user of private data is interested in more expressive queries  $q$ . The spectrum of such queries vastly varies. Examples include the mean value  $\frac{1}{n} \sum_{i=1}^n u_i$  of a collection of private data  $u_1, \dots, u_n$ , and solutions to optimization problems [8]. Our results are directly applicable to a broad family of queries which are approximated by private mechanisms built around the Laplace mechanism. Specifically, consider mechanisms based on the Laplace mechanism and have the form shown in Figure 5. The database of private data is initially preprocessed and, then, additive Laplace-distributed noise is used. The result is post-processed in order to maximize the accuracy of the response. Informally stated:

**Corollary 13.** *Let  $(\mathcal{U}, d)$  be a metric space of sensitive data,  $\mathcal{Y}$  be a set of responses, and  $\epsilon > 0$  be a privacy level. Let*

- $F : \mathcal{U} \rightarrow \Delta(\mathbb{R}^n)$  be a preprocessing step with sensitivity  $\beta$  that is invariant of  $\epsilon$ ,
- $\mathcal{L}_\epsilon : \mathbb{R}^n \rightarrow \Delta(\mathbb{R}^n)$  be the Laplace mechanism with parameter  $\epsilon$ :

$$\mathcal{L}_\epsilon u = u + V, \text{ where } V \sim e^{-\frac{\epsilon}{\beta} \|V\|_1}, \quad (40)$$

- $G_\epsilon : \mathbb{R}^n \rightarrow \Delta(\mathcal{Y})$  be a post-processing step.

Consider the  $\epsilon$ -private mechanism

$$G \circ \mathcal{L} \circ F : \mathcal{U} \rightarrow \Delta(\mathcal{Y}). \quad (41)$$

Then, there exists a composite mechanism that performs gradual release of sensitive data  $u \in \mathcal{U}$ .

Thus, our results are directly applicable to a existing privacy-aware mechanisms in e.g. smart grids [19], [2], and user's reports [2]. On the other hand, applying our results is not yet possible for mechanisms that do not fulfill this assumption, such as privately solving optimization problems with stochastic gradient descent [8].

In particular, Google's RAPPOR [2] is a mechanism that collects private data from multiple users for "crowdsourcing statistics" and can be expressed in terms of the Laplace mechanism. RAPPOR collects personal information from users such as the software features they use and the URLs they visited, and provides statistics of this information over a population of users. Algorithmically, a Bloom filter  $B$  is applied of size  $k$  is applied to each user's private data  $u$ :

$$B : \mathcal{U} \rightarrow \{0, 1\}^k, \quad y = [y_1, \dots, y_k] = B(u), \quad (42)$$

where  $\mathcal{U}$  is the space of private data, in particular, the set of all strings. Next, each bit  $y_i$  is perturbed with probability  $f$  and the result is memoised:

$$f : \{0, 1\}^k \rightarrow \{0, 1\}^k, \quad z = [z_1, \dots, z_k] = f(y), \text{ where } z_i = \begin{cases} 0, & \text{w.p. } \frac{1}{2}\alpha, \\ 1, & \text{w.p. } \frac{1}{2}\alpha, \\ y_i, & \text{w.p. } 1 - c_1, \end{cases} \quad (43)$$

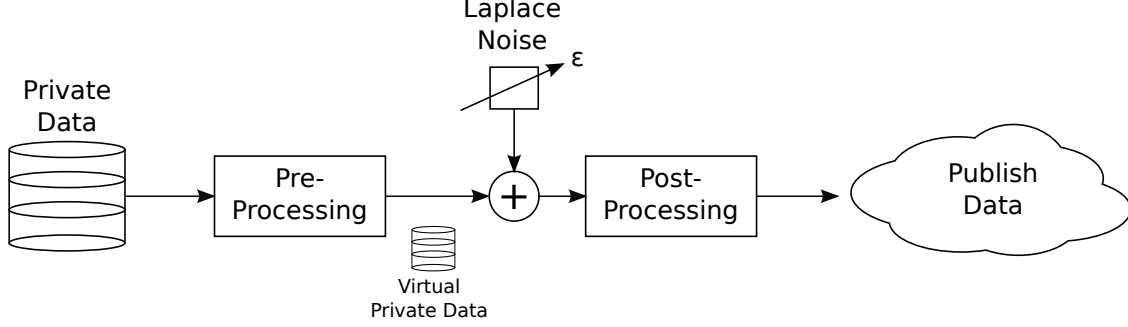


Figure 5: User 1 wants to share his sensitive data, such as his date of birth, in the a social network. Although, user 1 has no privacy concerns when sharing this information with his close friends 2 and 3, he has gradually increasing privacy issues for other members of the network. Specifically, a group  $A$  of distant users should not be able to collude and extract more information than what it is intended.

where “w.p.” stands for “with probability” and  $\alpha \in [0, 1]$  is a parameter. Finally, RAPPOR applies another perturbation each time a report is communicated to the server. This perturbation is equivalent to the map (43) but differently parametrized:

$$g : \{0, 1\}^k \rightarrow \{0, 1\}^k, \quad w = [w_1, \dots, w_k] = f(z), \quad \text{where } \mathbb{P}(w_i = 1) = \begin{cases} \beta, & \text{if } z_i = 1, \\ \gamma, & \text{if } z_i = 0, \end{cases} \quad (44)$$

where  $\beta, \gamma \in [0, 1]$  are parameters. RAPPOR’s differential privacy guarantees relax (increased  $\epsilon$ ) for small values of  $\alpha$  and  $\gamma$ , and large values of  $\beta$ .

An important limitation of RAPPOR is that parameters  $\alpha$ ,  $\beta$ , and  $\gamma$  are forever fixed. However, there are reasons that require the ability to update these values in a way that the privacy is relaxed and the accuracy is increased:

- Due to the non-trivial algorithm of decoding the reports, a tight accuracy-analysis is not possible. Instead, the accuracy of the system is evaluated once the system is bootstrapped.<sup>1</sup> Our results makes it possible to initialize the parameters with tight values  $\alpha \rightarrow 1$ ,  $\beta \rightarrow .5$ ,  $\gamma \rightarrow .5$ , and subsequently relax the parameters until a desired accuracy is achieved.
- Once a process or URL is suspected as malicious, the server would be interested in relaxing the privacy level and performing more accurate analysis of the potential threat. Once such a threat is identified, our result allows users to gradually relax their privacy parameters and the server can more confidently explore the potential threat.

In order to apply Theorems 10 and 12 to RAPPOR, we express the randomized maps (43) and (44) using the Laplace mechanism. Specifically, consider the functions  $\bar{f}$  and  $\bar{g}$  that add Laplace noise and project the result to  $\{0, 1\}$ :

$$\bar{f}(\psi) = \left[ \psi + V_f > \frac{1}{2} \right], \quad \text{where } V_f \sim \text{Lap} \left( \frac{1}{-2 \ln \alpha} \right), \quad (45)$$

$$\bar{g}(\zeta) = \left[ \zeta + V_g > \frac{\ln(2\gamma)}{\ln(4\gamma(1-\beta))} \right], \quad \text{where } V_g \sim \text{Lap} \left( \frac{1}{-\ln(4\beta(1-\gamma))} \right), \quad (46)$$

<sup>1</sup>Even in that case, estimating the actual accuracy can be challenging since it should be performed in a differential private way.

where  $\psi, \zeta \in \{0, 1\}$ ,  $\text{Lap}(b)$  is the Laplace distribution with parameter  $b$ , and  $[\text{expr}] \in \{0, 1\}$  is 1 if, and only if,  $\text{expr}$  is true. Note that functions  $\bar{f}$  and  $\bar{g}$  have the structure of Figure 5. Moreover, it can be shown that  $\bar{f}$  and  $\bar{g}$  applied component-wise to  $y$  and  $z$  are reformulations of the maps  $f$  and  $g$ . Therefore, privacy level relaxation is achieved by sampling noises  $V_f$  and  $V_g$  as suggested by our results.

## 4.2 Privacy in Social Networks

The context of social networks provides another setting where gradually releasing private data is critical. Consider a social network as a graph  $G = (V, E)$ , where  $V$  is the set of users and  $E$  the set of friendships between them. Each user owns a set of sensitive data that can include the date of birth, the number of friends and the city he currently resides. In the realm of social media, user's privacy concerns scale with the distance to other users of the network. Specifically, an individual is willing to share his private data with his close friends without any privacy guarantees, is skeptical about sharing this information with friends of his friends, and is alarmed to release his sensitive data over the entire social network. Therefore, an individual  $i$  chooses a different privacy level  $\epsilon_j$  for each user  $j \in V$  as a decreasing function between users  $i$  and  $j$ :

$$\epsilon_j = \frac{1}{d(i, j)}, \quad (47)$$

where  $d$  is a distance measure, e.g. the length of the shortest path between nodes  $i$  and  $j$ . Then, user  $i$  could generate an  $\epsilon_j$ -private response  $y_j$  independently for each member  $j$  of the network. However, more private information than desired is leaked. Specifically, consider the part of the social network shown in Figure 6, where user  $i = 1$  wishes to share his sensitive data  $u$ , such as her date of birth. Then, consider a group  $A \subseteq V$  of users residing far away from user 1 such that the privacy budget  $\epsilon_j$  allocated by user  $i$  to each member  $j$  of the group  $A$  is small:

$$d(1, j) \gg 1 \Rightarrow \epsilon_j \ll 1.$$

In the case that members of the large group  $A$  decide to collude, they can infer more information about the sensitive data  $u$ . Specifically, if a *large* group  $A$  averages the received responses  $\{y_j : j \in A\}$ , the exact value of sensitive data  $u$  is recovered. Indeed, composition theorem implies that only  $\left(\sum_{j \in A} \epsilon_j\right)$ -privacy of sensitive data  $u$  is guaranteed. For a large group  $A$ , this privacy level becomes very loose.

Our approach mitigates this issue. We assume that noisy versions of the private data are correlated and we design a mechanism that retains strong privacy guarantees. For real-valued sensitive data  $u$ , user 1 samples  $\{v_\epsilon : \epsilon > 0\}$  from the stochastic process  $\{V_\epsilon : \epsilon > 0\}$ , and responds to user  $j$  with  $y_j = u + v_{\epsilon_j}$ , as shown in Figure 7. In the case that a large group  $A$  of users colludes, they are unable to extract much more information. Specifically, such a collusion renders individual's sensitive information at most  $(\max_{j \in A} \epsilon_j)$ -differential private. This privacy budget is significantly tighter than the one derived in the naive application of differential privacy and corresponds to the best information that a member of the group  $A$  has. After all, if a close friend leaks sensitive information, it is impossible to revoke it.

## 5 Open Problems

Finally, we conjecture that gradually releasing private data can be extended to any query and is, therefore, an intrinsic property of differential privacy. This conjecture is a key ingredient for the




Figure 6: User 1 wants to share his sensitive data, such as his date of birth, in the a social network. Although, user 1 has no privacy concerns when sharing this information with his close friends 2 and 3, he has gradually increasing privacy issues for other members of the network. Specifically, a group  $A$  of distant users should not be able to collude and extract more information than what it is intended.




Figure 7: User 1 draws a single sample from the stochastic process  $\{V_\epsilon\}_{\epsilon>0}$  and responds to user  $i$  with  $y_i = u + V_{\epsilon_i}$ , where  $\epsilon_i$  is the privacy level against user  $i$ . Eventually, having access to more responses  $\{y_i\}_{i \in A}$  does not reveal more information about private data  $u$  than the best response  $\max_{i \in A} \epsilon_i$ .

existence of a frictionless market of private data. In such a market, owners of private data can gradually agree to a rational choice of privacy level. Moreover, buying the exact private data is expected to be extremely costly. Instead, people may choose to buy private data in “chunks”, in the sense of increasing privacy budgets. We conjecture that gradually releasing sensitive data without loss in accuracy is feasible for a broader family of privacy-preserving mechanisms beyond mechanisms that approximate identity queries. This work was focused mechanisms which are defined on real space or sensitive data  $\mathcal{U} = \mathbb{R}^n$  under an  $\ell_1$ -norm adjacency relation, and approximate the identity query.

## Acknowledgement

The authors would like to thank Aaron Roth for providing useful feedback and suggesting the application of our results to Google’s RAPPOR project.

## References

- [1] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *IEEE Symposium on Foundations of Computer Science*, 2007.
- [2] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1054–1067. ACM, 2014.
- [3] Moritz Hardt and Kunal Talwar. On the geometry of differential privacy. In *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 705–714. ACM, 2010.
- [4] Chao Li, Michael Hay, Vibhor Rastogi, Gerome Miklau, and Andrew McGregor. Optimizing linear counting queries under differential privacy. In *Proceedings of the twenty-ninth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 123–134. ACM, 2010.
- [5] Jonathan Ullman. Answering  $n \{2+ o(1)\}$  counting queries with differential privacy is hard. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 361–370. ACM, 2013.
- [6] Jerome Le Ny and George J Pappas. Differentially private filtering. *Automatic Control, IEEE Transactions on*, 59(2):341–354, 2014.
- [7] Anupam Gupta, Katrina Ligett, Frank McSherry, Aaron Roth, and Kunal Talwar. Differentially private combinatorial optimization. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1106–1125. Society for Industrial and Applied Mathematics, 2010.
- [8] Shuo Han, Ufuk Topcu, and George J Pappas. Differentially private distributed constrained optimization. *arXiv preprint arXiv:1411.4105*, 2014.
- [9] Justin Hsu, Zhiyi Huang, Aaron Roth, Tim Roughgarden, and Zhiwei Steven Wu. Private matchings and allocations. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 21–30. ACM, 2014.

- [10] Jason Reed and Benjamin C Pierce. Distance makes the types grow stronger: a calculus for differential privacy. In *ACM Sigplan Notices*, 2010.
- [11] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM symposium on Theory of computing*, pages 715–724. ACM, 2010.
- [12] Konstantinos Chatzikokolakis, Miguel E Andrés, Nicolás Emilio Bordenabe, and Catuscia Palamidessi. Broadening the scope of differential privacy using metrics. In *Privacy Enhancing Technologies*, pages 82–102. Springer, 2013.
- [13] Fragkiskos Koufogiannis, Shuo Han, and George Pappas. Optimality of the laplace mechanism in differential privacy. *arXiv preprint arXiv:1504.00065*, 2015.
- [14] Yu Wang, Zhenqi Huang, Sayan Mitra, and Geir E Dullerud. Entropy-minimizing mechanism for differential privacy of discrete-time linear feedback systems. In *IEEE Conference on Decision and Control*, 2014.
- [15] Xiaokui Xiao, Gabriel Bender, Michael Hay, and Johannes Gehrke. ireduct: Differential privacy with reduced relative errors. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*, pages 229–240. ACM, 2011.
- [16] Cynthia Dwork. Differential privacy. In *Automata, languages and programming*, 2006.
- [17] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Theoretical Computer Science*, 9(3-4):211–407, 2013.
- [18] Quan Geng and Pramod Viswanath. The optimal mechanism in differential privacy. *arXiv preprint arXiv:1212.1186*, 2012.
- [19] Fragkiskos Koufogiannis, Shuo Han, and George J Pappas. Computation of privacy-preserving prices in smart grids. In *IEEE Conference on Decision and Control*, 2014.